The High Bandwidth, 256KB 2nd Level Cache on an Itanium™ Microprocessor.

Tom Grutkowski

Reid Riedlinger

Intel Corporation and Hewlett Packard

McKinley L2 Features

- Hub for memory operations in McKinley memory sub-system, enforcing all memory ordering requirements.
- Non-blocking/out-of-order address queue.
- Handles all L1 Instruction and L1 Data Misses, along with all stores.
- All Floating Point Memory Access are served directly by L2, allowing for integer optimization of L1D.
- High sustainable bandwidth to core:
  - 64GB/S to L1 Caches @1Ghz (perfectly interleaved L1I and L1D FILL requests)
  - 40 GB/S to FPU (4 ext. precision operands/clock from independent addresses)
- Implements all architecturally defined semaphore instructions.

The L2 within McKinley

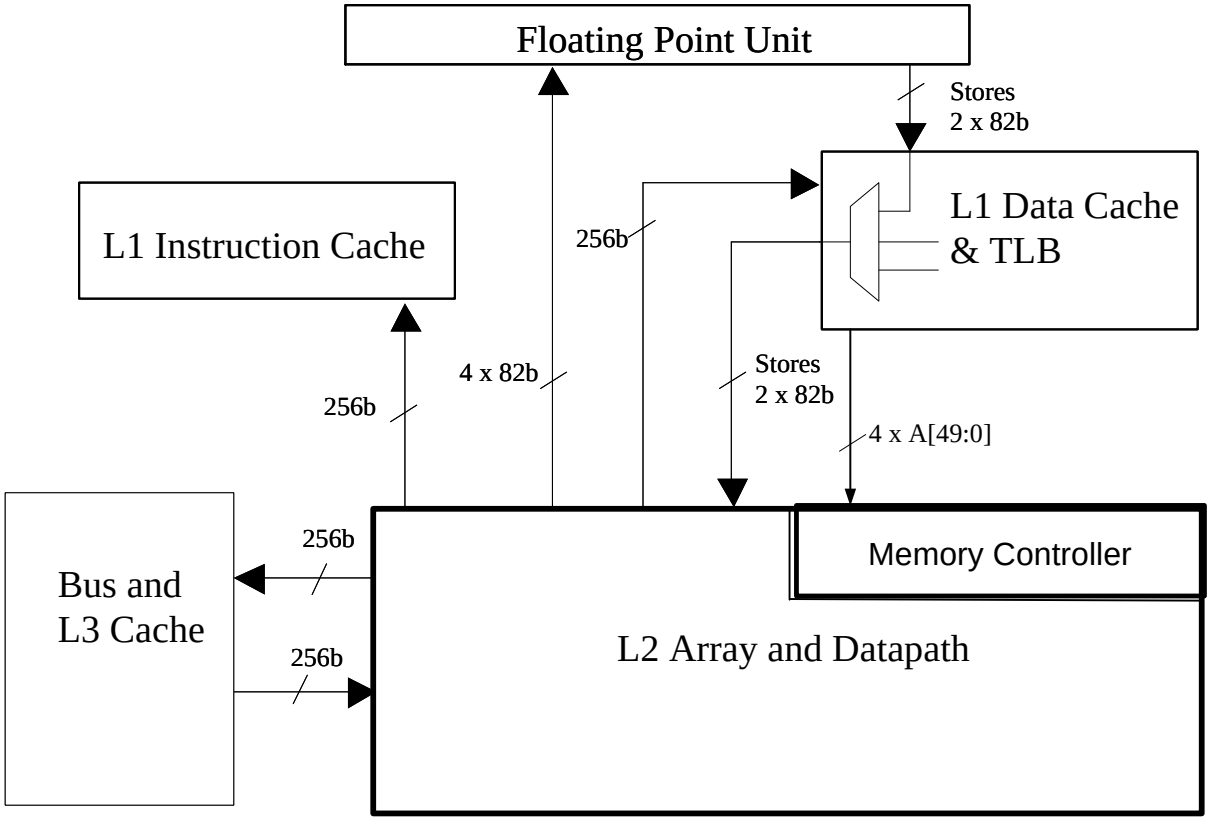

Inside the L2

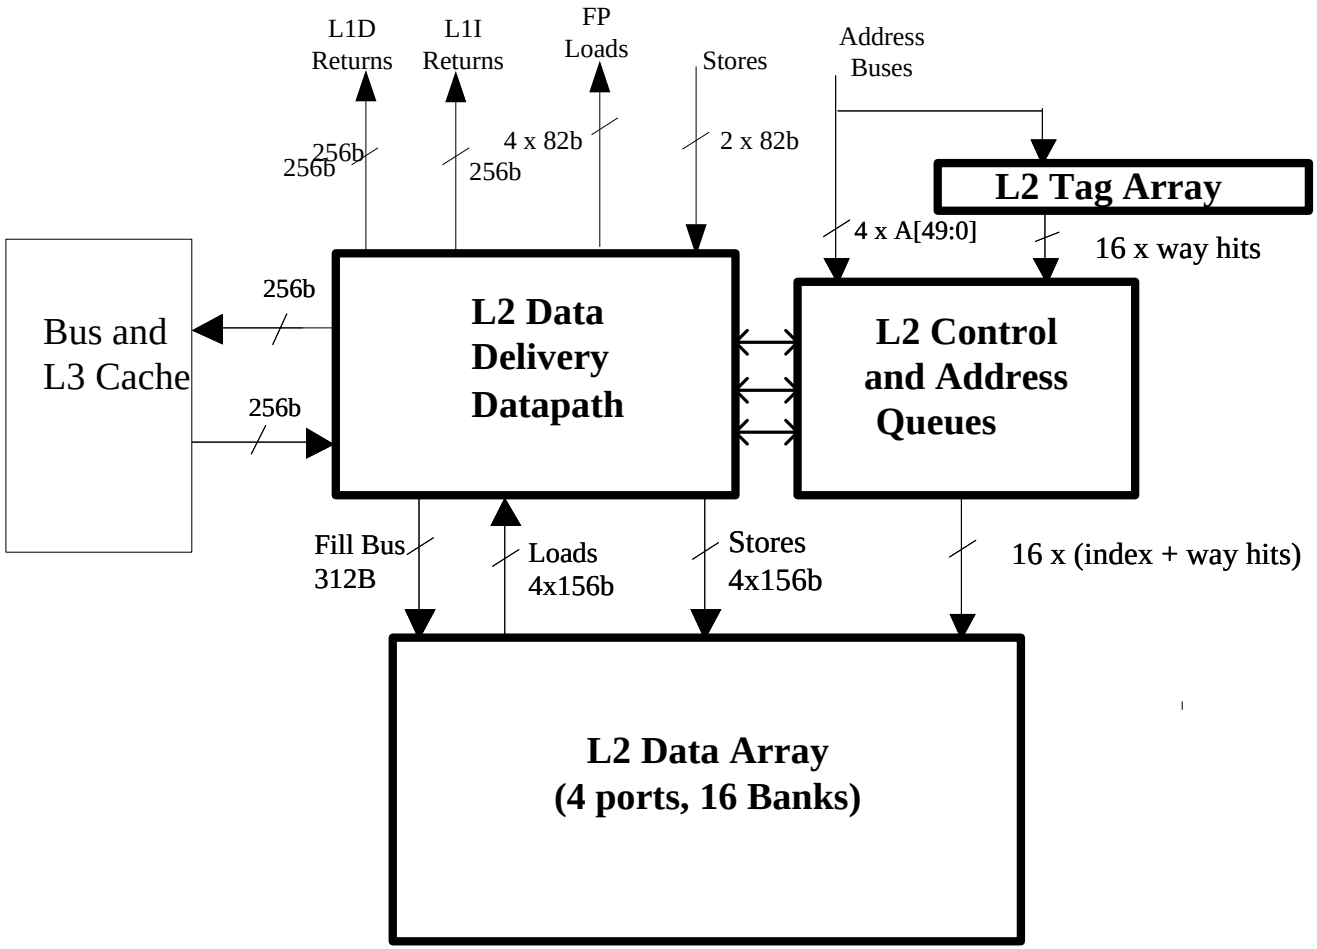

Die Photo

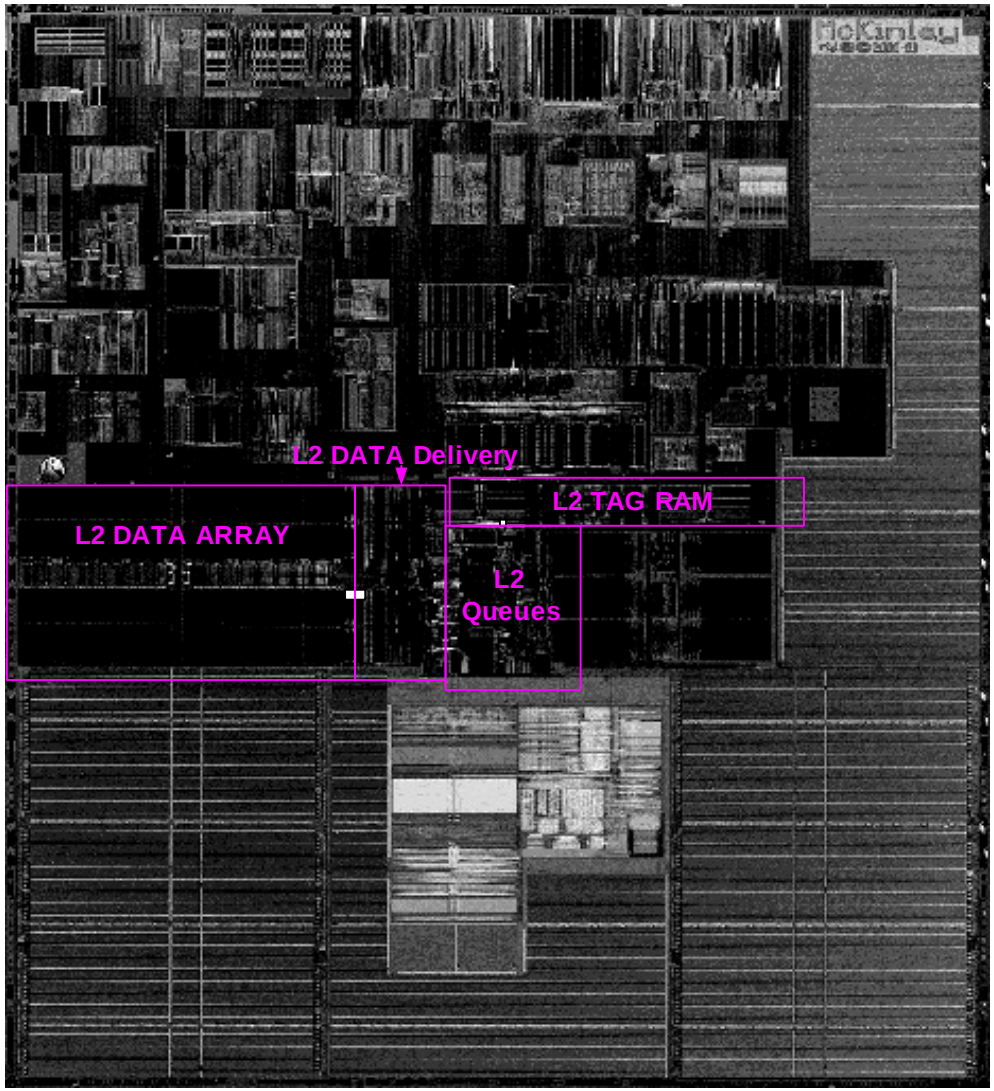

L2 has it own pipeline.

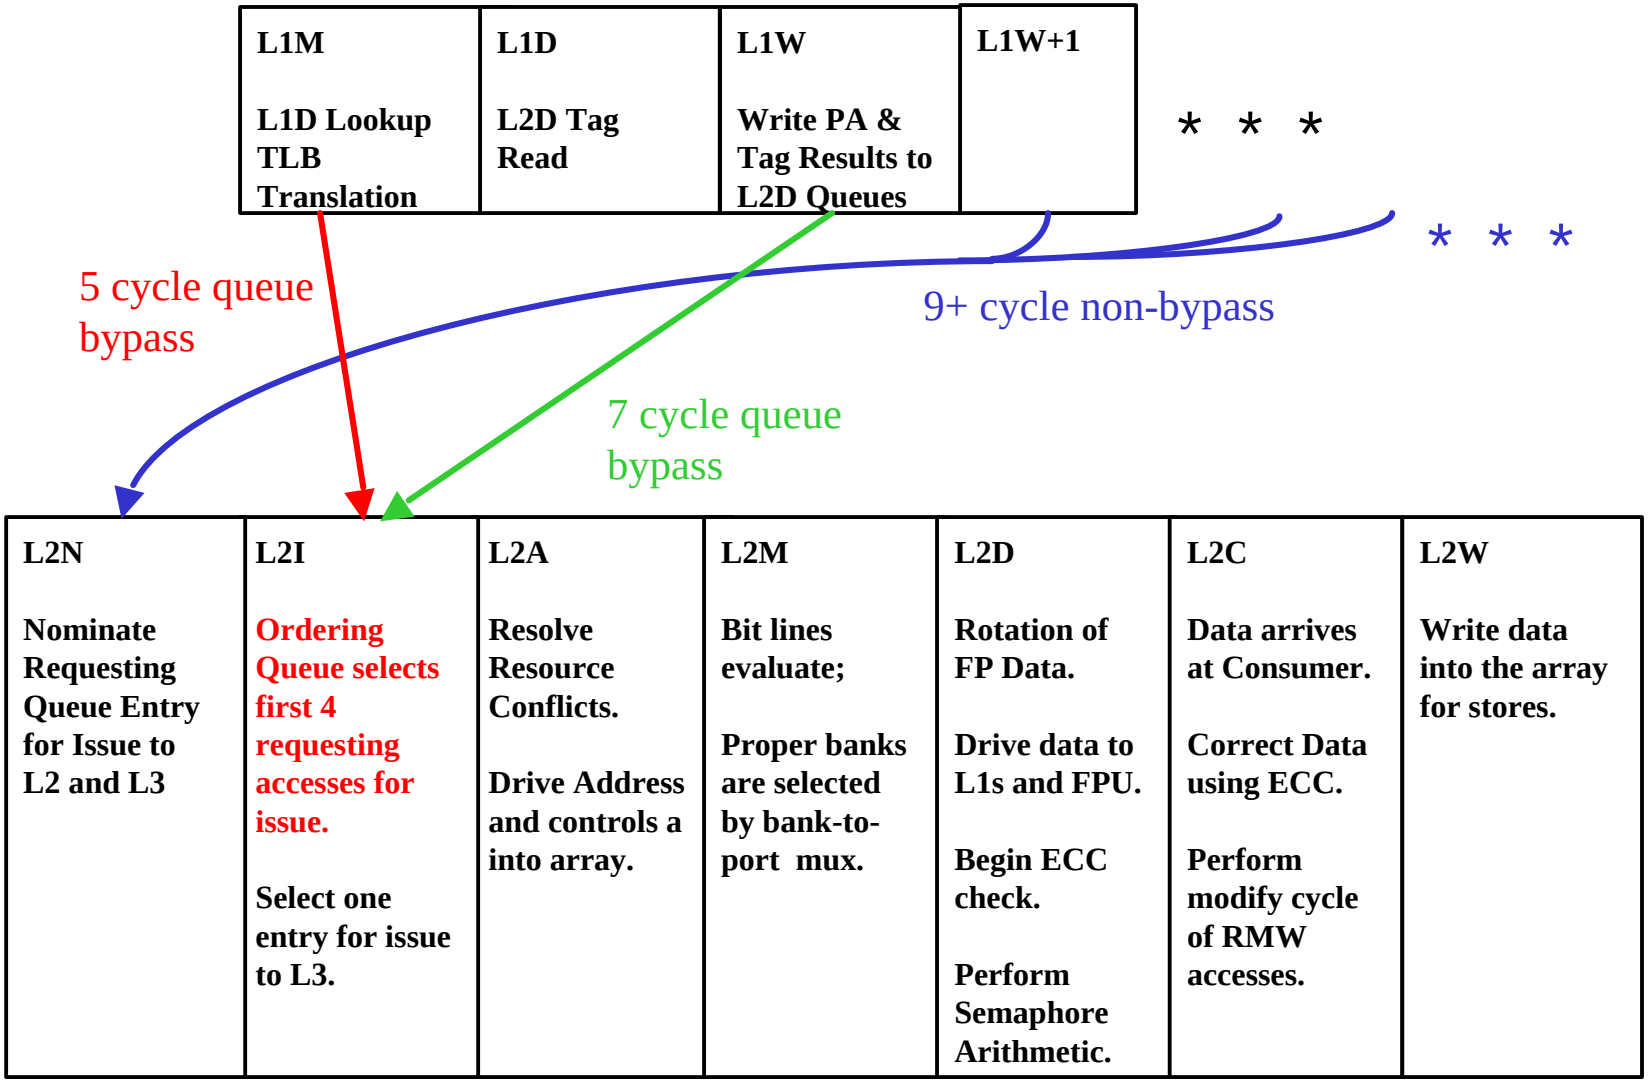

A look at the 5 Cycle Bypass Case

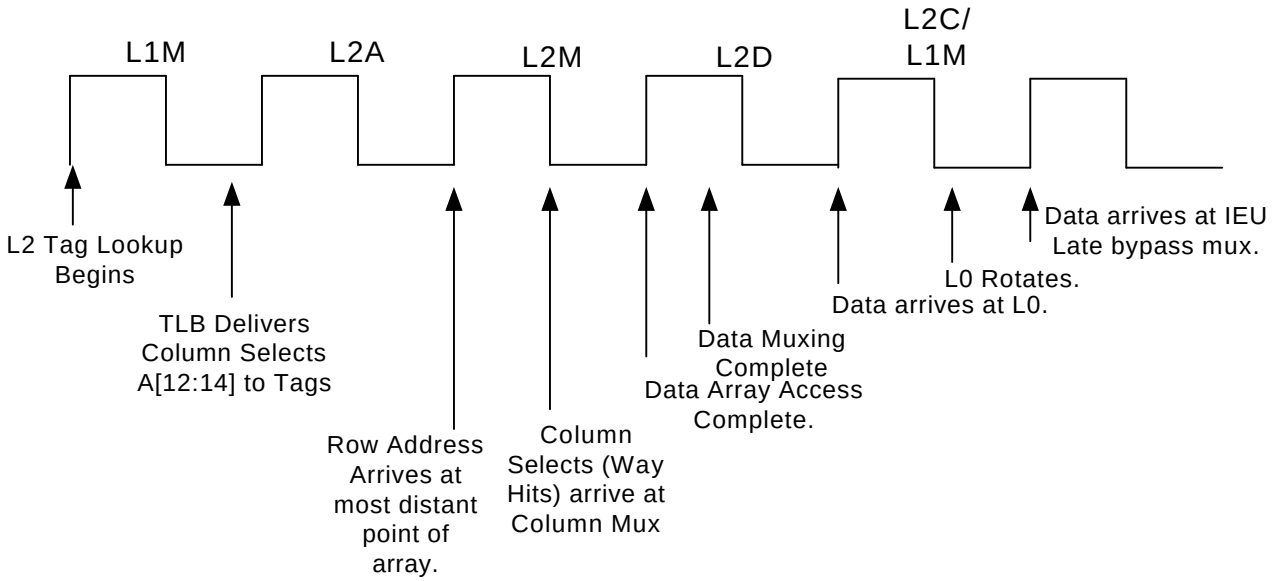

- Benefits of data bypassing.
  - 7.6% (FP) and 5.8 (Int) from Data Bypassing
  - 0.9% (FP) and 2.2 (Int) from Instruction Bypassing.

L2 Tag Access

- L2's 4-ported Tag Ram
  - Implemented using two identical 2-ported tag arrays
  - Uses cell similar to that used in McKinley's dual ported L1 array. (6T single ended write/ dual ended read)

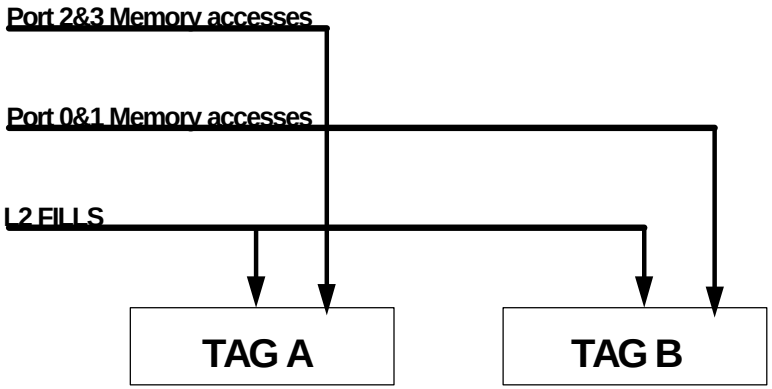

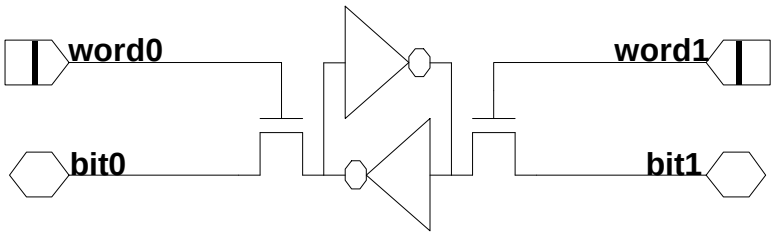

L2 Array Features

- 256K Byte Data Array
  - 4 ported for loads.
    - 16 Bytes per port are read each clock
  - Protected by ECC
    - 32 Bit ECC Word
    - Single Bit Error Correction/Double Bit Detection
  - Stores are performed with 32 bit granularity without need for RMW cycle (very helpful for ia32 support).
- Capable of simultaneously performing:
  - 4 Memory Loads from independent addresses
  - 4 Stores to independent address
  - 128 B single clock fill

L2 Array Access

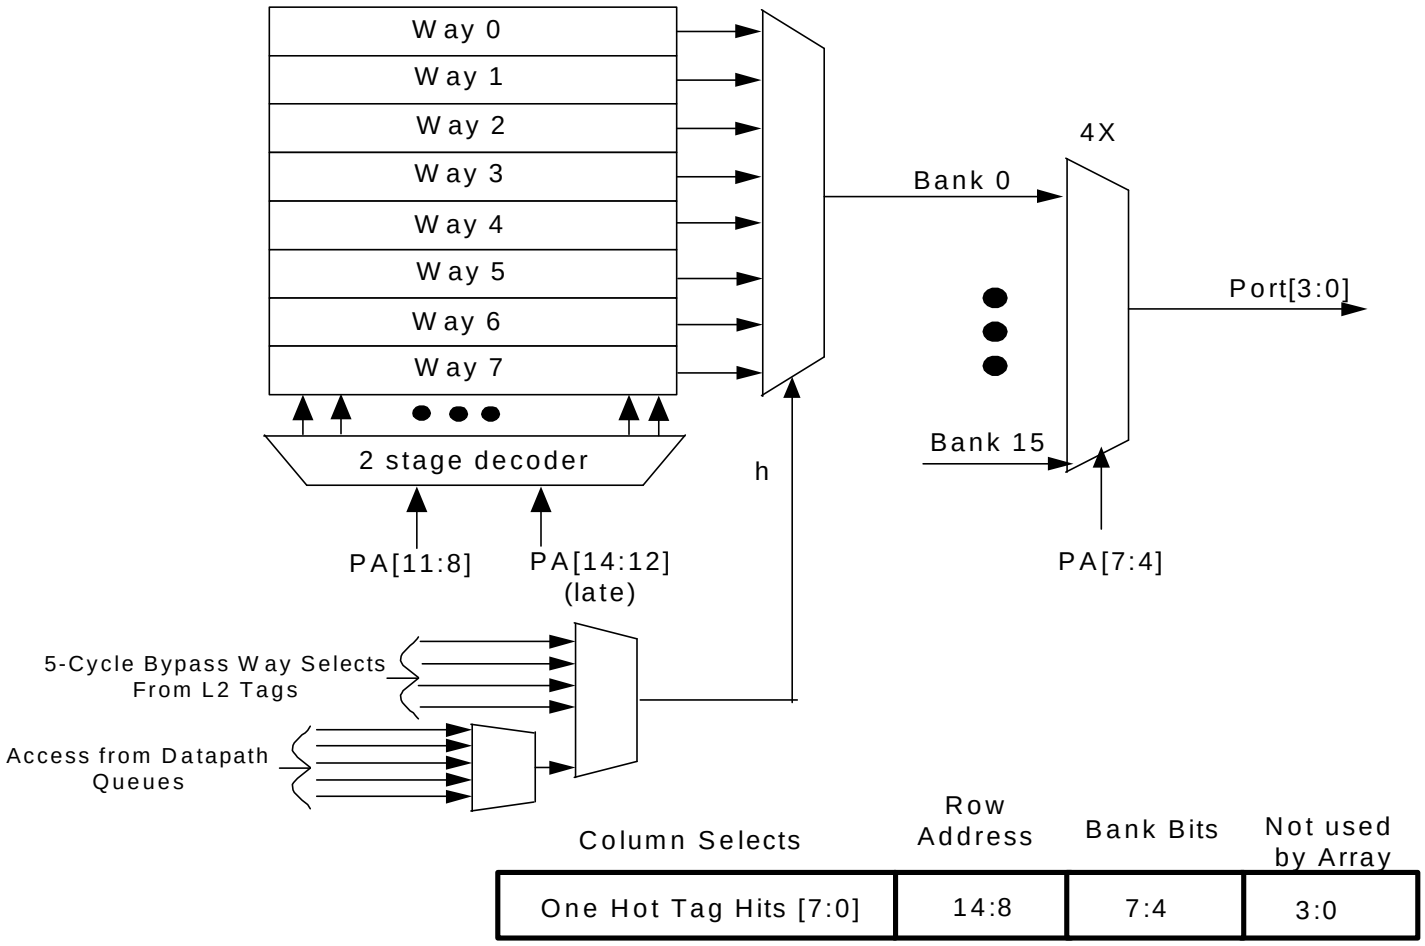

L2 Array Circuitry

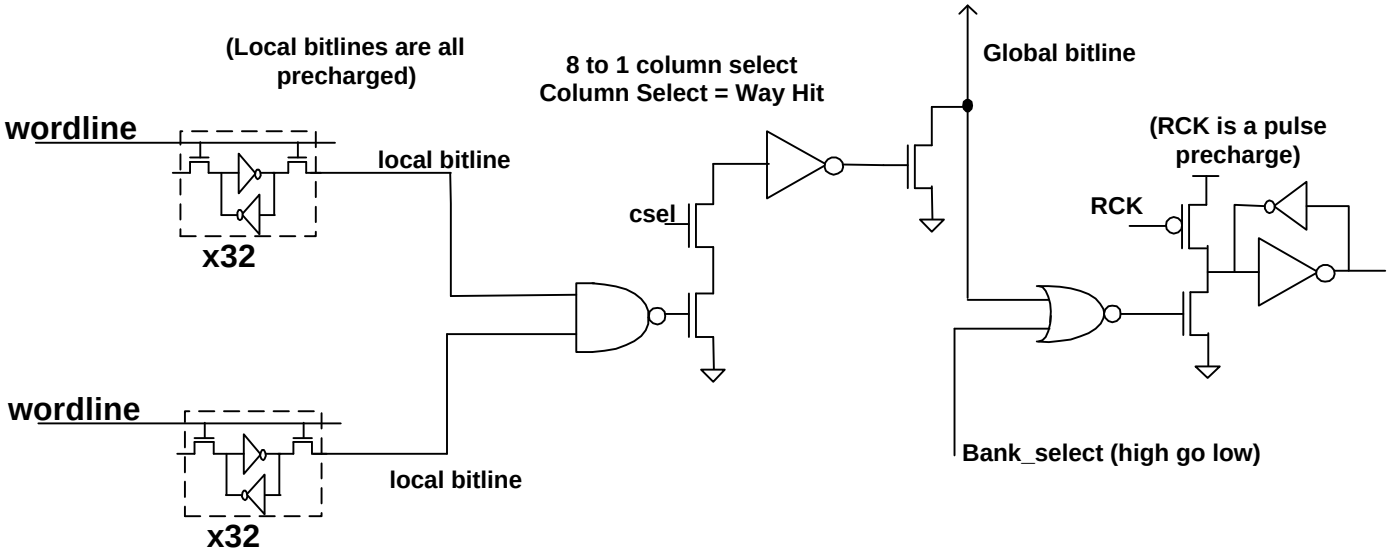

L2 Queues

- L2 Queue blocks
  - Forms the central clearing house for all address transactions in McKinley memory sub system.
  - 15 individual subunits.
  - 1.4 million transistors, and nearly 6000 top level routes.
- Main Queue Structure holds outstanding L2 Data accesses and L3/Bus access not yet accepted by L3/Bus
  - Main Queue Depth is 32 entries.
  - Select four L2 accesses per clock for issue to L2 array.
    - Ideally these accesses have no bank conflicts.
    - Selected accesses should obey all software defined ordering constraints.
  - Queue must also select a single access each clock for issue to L3/Bus.

Nominate and Issue Sequence

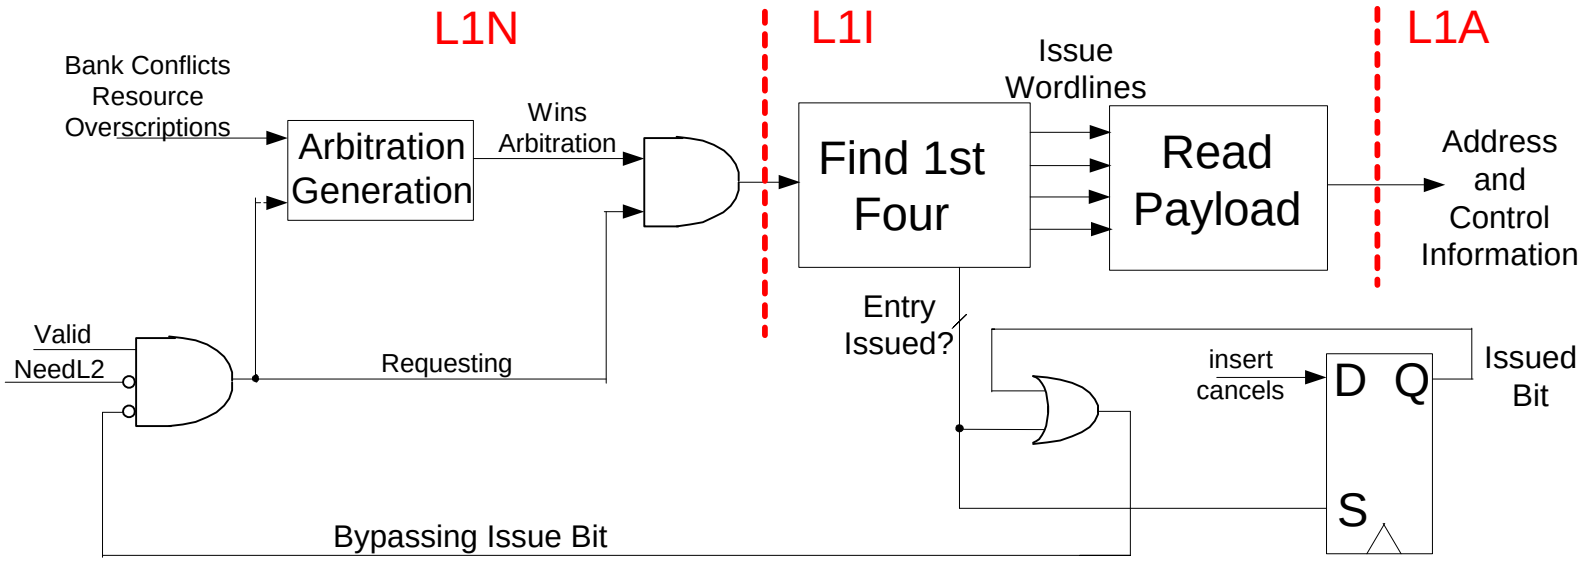

Techniques to speed find first four circuitry.

- Because entries remain in queue for 4 clocks following issue,
  - Jump over these already issued accesses in L1N, and
  - Only reset L1I search when cancels/ECC errors occur.
  - This limits the needed FF4 search depth to only 16 entries,
  - And provides for a scalable solution as queue size increases.
- Wrap queue in an interleaved manner.
- Use 4-wide ‘look-ahead’ technique to implement find first circuitry.

Illustration of Jump Start

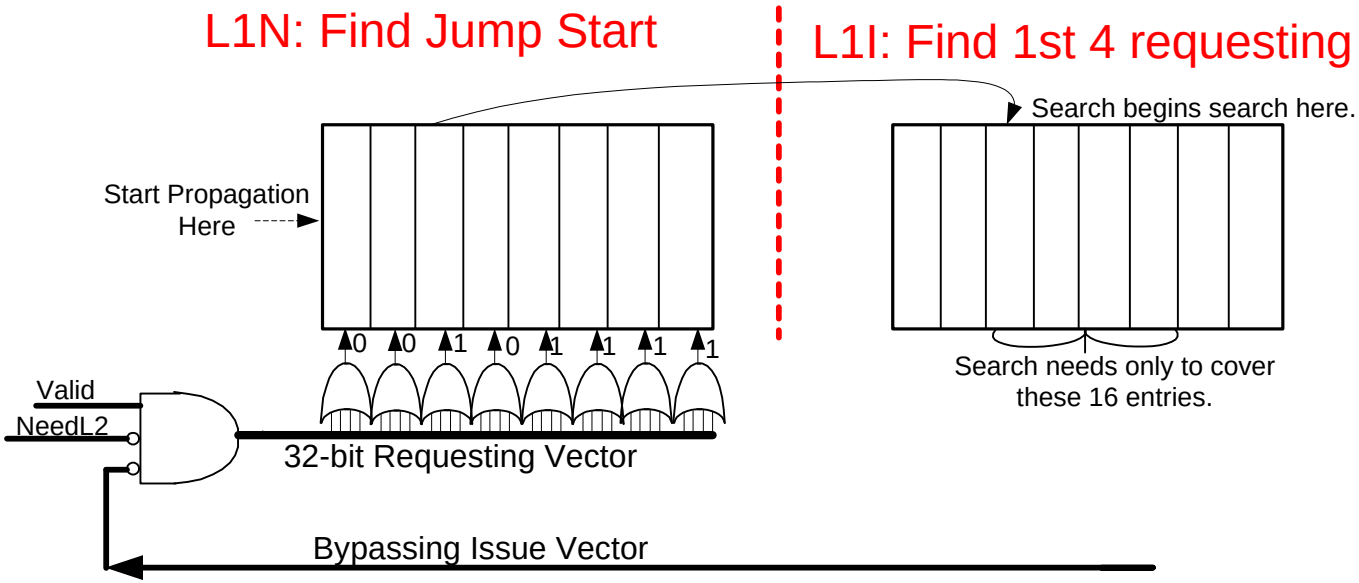

Interleaved Groups of 4

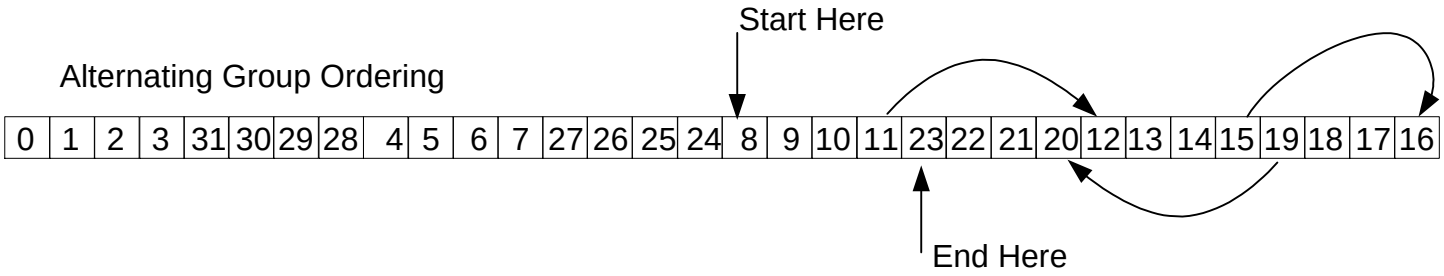

- No long wrap around case, therefore all entries can be done with identical circuitry.
- Makes design very scalable. Timing of search remains constant regardless of number of entries.

Find FF4 Group Circuitry

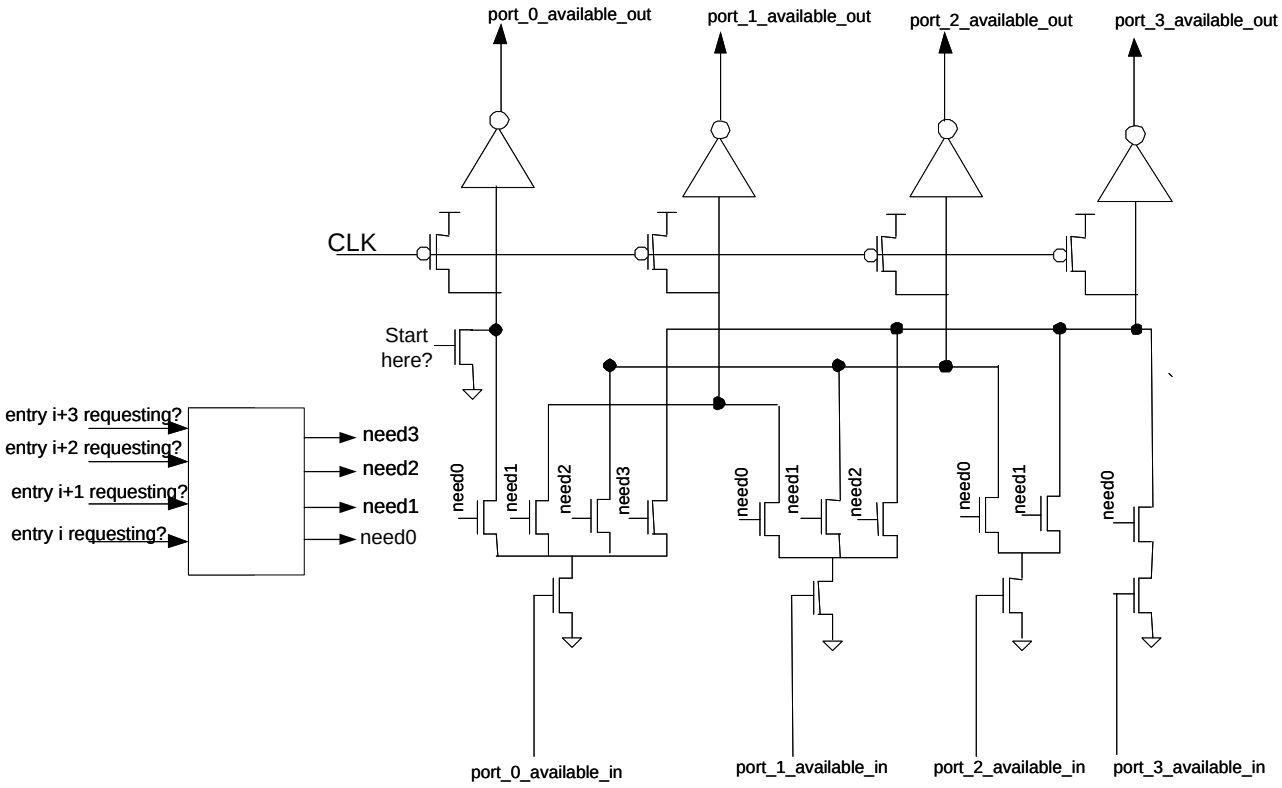

Summary

- Aggressive circuit and floor planning techniques enabled the L2 queues to fall within 8.6mm^2 area.
- Implemented 5 and 7 cycle bypassing.
  - ~8% performance boost on FP/INT code
- High bandwidth design ensures that FPU and IEU are not data starved.
- Took on all FP memory operations and memory ordering enforcement, thereby helping the L1D to achieve 1 cycle operation.